Understanding the structure of volatility risks

George Panayotov *
University of Maryland, College Park

March 15, 2005

Abstract

Studies of asset returns time-series provide strong evidence that at least two stochastic factors drive volatility. This paper investigates whether two volatility risks are priced in the stock option market and estimates volatility risk prices in a cross-section of stock option returns. The paper finds that the risk of changes in short-term volatility is significantly negatively priced, which agrees with previous studies of the pricing of a single volatility risk. The paper finds also that a second volatility risk, embedded in longer-term volatility is significantly positively priced. The difference in the pricing of short- and long-term volatility risks is economically significant - option combinations allowing investors to sell short-term volatility and buy long-term volatility offer average profits up to 20% per month.

*I am grateful to Dilip Madan for all the guidance and support. I thank Doron Avramov, Gurdip Bakshi, Michael Fu, Steve Heston, Gerry Hoberg, Gergana Jostova, Ben Kedem, Mark Loewenstein and Oghenovo Obrimah for helpful comments. I especially thank Haluk Unal and Levent Guntay whose suggestions greatly improved the paper. Any errors are mine.

1 Introduction

Recent studies provide evidence that market volatility[1] risk is priced in the stock option market (e.g. Chernov and Ghysels (2000), Benzoni (2001), Coval and Shumway (2001), Bakshi and Kapadia (2003), Carr and Wu (2004)). These studies typically find a negative volatility risk price, suggesting that investors are ready to pay a premium for exposure to the risk of changes in volatility. All these studies consider the price of risk, embedded in a *single* volatility factor.

In contrast, time-series studies find that *more than one* stochastic factor drives asset returns volatility. Engle and Lee (1998) find support for a model with two volatility factors - permanent (trend) and transitory (mean-reverting towards the trend). Gallant, Hsu and Tauchen (1999), Alizadeh, Brandt and Diebold (2002) and Chernov, Gallant, Ghysels and Tauchen (2003) estimate models with one highly persistent and one quickly mean-reverting volatility factor and show that they dominate over one-factor specifications for volatility[2].

Motivated by the results of the time-series studies, this paper investigates whether the risks in two volatility risks are priced in the stock option market. I construct the volatility factors using implied volatilities from index options with different maturities (between one month and one year).

My main finding is that two volatility risks are indeed priced in the stock option market. The risk of changes in short-term volatility is significantly negatively priced. This result is consistent with the previous volatility risk pricing literature, which uses relatively short-term options and finds negative price of volatility risk. In addition, the paper reports a novel finding - I find that another risk, embedded in longer-term volatility is significantly positively priced. The positive risk price indicates that investors require positive compensation for exposure to long-term volatility risk.

This finding complements previous results on the pricing of two volatility factors in the stock market: Engle and Lee (1998) find that the permanent (or persistent) factor in volatility is significantly positively correlated with the market risk premium, while the transitory factor is not. MacKinlay and Park (2004) confirm the positive correlation of the permanent volatility factor with the risk premium and also find a time-varying and typically negative correlation of the transitory volatility factor with the risk premium.

The differential pricing of volatility risks in the stock option market, found in this paper, is also economically significant, as evidenced by returns on long calendar spreads[3].

[1] This paper only considers stock market volatility and does not touch upon the volatility of individual stocks. For brevity I will refer to market volatility as "volatility".

[2] See also Andersen and Bollerslev (1997), Liesenfeld (2001), Jones (2003).

[3] A long calendar spread is a combination of a short position in an option with short maturity and a long position in an option on the same name, of the same type and with the same strike, but of a longer maturity.

Expected returns on a calendar spread reflect mostly the compensations for volatility risks embedded in the two components of the spread. A short position in a negatively priced volatility risk (short-term) combined with a long position in a positively priced volatility risk (long-term) should then have a positive expected return. I calculate returns on calendar spreads written on a number of index and individual options and find that, in full support of the statistical estimations, spreads on puts gain an impressive 20% monthly on average, while spreads on calls gain about 12% on average. Transaction costs would reduce these numbers, but still, a pronounced difference in volatility risk prices can be captured using calendar spreads.

To perform the empirical tests I construct time-series of daily returns to options of several fixed levels of moneyness and maturity[4]. For this construction I use options on six stock indexes and twenty two individual stocks. I estimate volatility risk prices in the cross-section of expected option returns using the Fama-MacBeth approach and Generalized Method of Moments (GMM). These methods have not been applied to a cross-section of option returns in previous studies of volatility risk pricing. The cross-sectional analysis allows easily to decompose implied volatility and to estimate separately the prices of the risks in different components. This decomposition turns out to be essential for the disentangling of the risks, embedded in long-term implied volatility.

What is the economic interpretation of the different volatility risk prices found in the paper? There is still little theoretical work on the pricing of more than one volatility risks. Tauchen (2004) studies a model with two consumption-related stochastic volatility factors, which generates endogenously two-factor volatility of stock returns. One feature of Tauchen's model is that the risk prices of the two volatility factors are necessarily of the same sign. This paper offers an alternative model, which is able to generate volatility risk prices of different signs, consistent with the empirical findings. In this model the representative investor's utility function is concave in one source of risk and convex in a second source of risk (for reasonable levels of risk aversion). Such a utility function is closely related to multiplicative habit formation models (e.g. Abel (1990)). Both risk sources exhibit stochastic volatility. The negative volatility risk price is associated to concavity of the utility function, whereas the positively priced volatility risk is associated to convexity of the utility function.

The rest of the paper is organized as follows. Section 2 discusses the pricing of volatility risks in models with one and two volatility factors. Section 3 describes the construction of option returns and the volatility risk factors. Section 4 presents the estimation results and Section 5 concludes.

[4]Such constructs have been used before on a limited scale - e.g. the CBOE used to derive the price of an at-the-money 30-day option to calculate the Volatility Index (VIX) from 1986 to 1993; Buraschi and Jackwerth (2001) use 45-day index options with fixed moneyness levels close to at-the-money

2 Volatility risk prices in models with one and two volatility factors

This section considers first the pricing of volatility risk in a model with a single stochastic volatility factor. It complements previous empirical findings of a negative volatility risk price by deriving analytically such a negative price in a stochastic volatility model of the type studied in Heston (1993). A negative price of volatility risk is obtained in this model even if no correlation between asset returns and volatility is assumed. This model is later extended to include a second volatility factor. The predictions of the extended model are consistent with the empirical findings of this paper. Next, the section discusses the evidence for two volatility factors and specifies the relation that is tested empirically in the rest of the paper.

2.1 The price of a single volatility risk

Consider a standard economy with a single volatility risk. The representative investor in this economy holds the market portfolio and has power utility over the terminal value of this portfolio: $U_T = \frac{(S_T)^{1-\lambda}}{1-\lambda}$. The pricing kernel process in this economy is of the form:

$$\Lambda_t = E_t\left[S_T^{-\lambda}\right] \tag{2.1}$$

Expectation is taken under the statistical measure, S denotes the value of the market portfolio and λ is the risk aversion coefficient. Assume the following dynamics for S:

$$\frac{dS_t}{S_t} = D_t^S dt + \sqrt{\sigma_t^S} dW_t^S \tag{2.2}$$

where W_t^S is standard Brownian motion. The drift D_t^S is not modeled explicitly, since it does not affect the pricing kernel in the economy. The volatility $\sqrt{\sigma_t^S}$ is stochastic. Assume further that σ_t^S follows a CIR process, which is solution to the following stochastic differential equation:

$$d\sigma_t^S = k(\theta - \sigma_t^S)dt + \eta\sqrt{\sigma_t^S} dW_t \tag{2.3}$$

The model (2)-(3) is Heston's (1993) stochastic volatility model (with possibly time-varying drift D_t^S). The Brownian motions W_t^S and W_t are assumed here to be uncorrelated. I discuss below the implication for volatility risk pricing of the correlation between W_t^S and W_t.

Appendix A contains the proof of the following:

Proposition 1 *The stochastic discount factor ζ_t for the economy with one stochastic volatility factor described in (2.1)-(2.3) is given by*

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda \frac{dS_t}{S_t} + \lambda^\sigma d\sigma_t \tag{2.4}$$

where the price of market risk λ is strictly positive and the price of volatility risk λ^σ is strictly negative.

The economic intuition for this negative volatility risk price can be provided by the concavity of the utility function. Higher volatility results in lower expected utility (due to concavity). Then any asset, which is positively correlated with volatility has high payoff precisely when expected utility is low. Hence, such an asset acts as insurance and investors are ready to pay a premium for having it in their portfolio. This argument is well known, e.g. from the risk management literature.

Given the linear form of ζ_t, it is easy to test the model's prediction for the sign of λ^σ in cross-sectional regressions of the type

$$E\left[R^i\right] = \lambda^M \beta_i^M + \lambda^\sigma \beta_i^\sigma + \gamma^i \tag{2.5}$$

where: $E\left[R^i\right]$ is expected excess return on test asset i; the betas are obtained in time-series regressions of asset i's returns on proxies for market risk $\left(\frac{dS_t}{S_t}\right)$ and for volatility risk ($d\sigma_t$); λ^M and λ^σ are the prices of market risk and volatility risk respectively and γ^i is the pricing error. If the two risk factors are normalized to unit variance, this beta-pricing representation yields risk prices equivalent to those in the stochastic discount factor model (2.4) (e.g. Cochrane (2001), Ch. 6).

Ang et al. (2004) test a model similar to (2.5) and find significantly negative price of volatility risk in the cross-section of expected stock returns. Empirical tests of (2.5) involving option returns are reported in Section 4 in this paper and also support a negative price of volatility risk.

These tests of the simple relation (2.5) are consistent with most of the previous studies of volatility risk pricing in the options markets. Benzoni (2001), Pan (2002), Doran and Ronn (2003), among others, use parametric option-pricing models to estimate a negative volatility risk price from option prices and time series of stock market returns. Coval and Shumway (2001) argue that if volatility risk were not priced, then short delta-neutral at-the-money straddles should earn minus the risk free-rate. In contrast, they find 3% average gain per week, which is (tentatively) interpreted as evidence that market volatility risk is negatively priced. Within a general two-dimensional diffusion model for asset returns, Bakshi and Kapadia (2003) derive that expected returns to delta-hedged options are positive (negative) exactly when the price of volatility risk is positive (negative) and find significantly negative returns. Carr and Wu (2004) construct synthetic variance swap rates from option prices and compare them to realized variance - the variance risk premium obtained in this way is significantly negative.

The derivation of (2.4) assumes that the two Brownian motions are uncorrelated. If ρ is a non-zero correlation between the Brownian motions, then the volatility risk price has two components: $\rho\lambda$ and $\lambda^{\sigma}\sqrt{1-\rho^2}$. Previous empirical studies have often argued that the negative price of volatility risk they find is due to the negative correlation between changes in volatility and stock returns. So, they focus on the negative $\rho\lambda$ term. This is a powerful argument, given that this negative correlation is among the best-established stylized facts in empirical finance (e.g. Black (1976)). Exposure to volatility risk is thus seen as hedging against market downturns and the negative volatility risk price is seen as the premium investors pay for this hedge. However, focusing on the negative correlation leaves aside the second term in volatility risk price ($\lambda^{\sigma}\sqrt{1-\rho^2}$). That such a term can be important is indicated e.g. in Carr and Wu (2004) - they find that even after accounting for the correlation between market and volatility risks, there still remains a large unexplained negative component in expected returns to variance swaps. A negative λ^{σ} as derived above is consistent with this finding. To explore the relative significance of the different components of volatility risk price, this paper reports empirical tests which include a market risk factor and uncorrelated volatility risks.

2.2 Models with multiple volatility factors

The study of models with multiple volatility factors has been provoked partly by the observation that the volatility of lower frequency returns is more persistent than the volatility of higher frequency returns. This pattern can be explained by the presence of more than one volatility factors, each with a different level of persistence. Such factors have been interpreted in several ways in the literature:

Andersen and Bollerslev (1997) argue that volatility is driven by heterogeneous information arrival processes with different persistence. Sudden bursts of volatility are typically dominated by the less persistent processes, which die out as time passes to make the more persistent processes influential. Muller et al. (1997) focus on heterogeneous agents, rather than on heterogeneous information processes. They argue that different market agents have different time horizons. The short-term investors evaluate the market more often and perceive the long-term persistent changes in volatility as changes in the average level of volatility at their time scale; in turn, long-term traders perceive short-term changes as random fluctuations around a trend. Liesenfeld (2001) argues that investors' sensitivity to new information is not constant but time-varying and is thus a separate source of randomness in the economy. He finds that the short-term movements of volatility are primarily driven by the information arrival process, while the long-term movements are driven by the sensitivity to news.

MacKinlay and Park (2004) study the correlations between the expected market risk premium and two components of volatility - permanent and transitory. The permanent component is highly persistent and is significantly positively priced in the risk premium, suggesting a positive risk-returns relation. The transitory component is highly volatile and tends to be negatively priced in the risk premium. This component is related to extreme market movements, transitory market regulations, etc. which can be dominating volatility dynamics over certain periods of time.

Tauchen (2004) is the first study to incorporate two stochastic volatility factors in a general equilibrium framework[5]. The two-factor volatility structure is introduced by assuming consumption growth with stochastic volatility, whereby the volatility process itself exhibits stochastic volatility (this is the second source of randomness in volatility). The model generates endogenously a two-factor conditional volatility of the stock return process. It also generates a negative correlation between stock returns and their conditional volatility as observed empirically in data. One feature of his model is that the risk premia on the two volatility factors are both multiples of the same stochastic process (the volatility of consumption volatility) and are necessarily of the same sign. The model thus imposes a restriction on the possible values of volatility risk prices.

The model with two volatility factors which is tested in this paper does not impose *a'priori* restrictions on volatility risk prices. In analogy with (2.5) I estimate cross-sectional regressions of the type:

$$E\left[R^i\right] = \lambda^M \beta_i^M + \lambda^L \beta_i^S + \lambda^L \beta_i^S + \gamma^i. \tag{2.6}$$

where:

- the test assets are options (unhedged and delta-hedged) on a number of stock indexes and individual stocks and $E\left[R^i\right]$ denotes expected excess returns on option i;

- the betas are obtained in time-series regressions of option i's returns on proxies for market risk $\left(\frac{dS_t}{S_t}\right)$ and for two volatility risks $\left(d\sigma_t^S \text{ and } d\sigma_t^L\right)$; in particular, I consider a short- and a and long-term volatility, denoted by superscripts S and L resp.

- λ^M, λ^S and λ^L are the prices of market risk and the two volatility risks respectively, and γ^i is the pricing error.

Equation (2.6) presents a three-factor model with a linear stochastic discount factor. Similar specifications have been widely and successfully employed in empirical asset-pricing tests. It would be interesting, though, to search for an economic justification for equation (2.6). One possibility would be to formally extend the model in (2.1)-(2.3) by

[5] In a related work Bansal and Yaron (2004) model consumption growth as containing a persistent predictable component plus noise. Stochastic volatility is incorporated both in the persistent component and the noise. However, only one source of randomness in volatility is assumed in their model, common to both components of consumption growth.

adding a second source of randomness with stochastic volatility to (2.2). However, this approach would come up with the prediction that both volatility risk prices are negative. An alternative economic model consistent with (2.6) is presented in Appendix B. This model is related to habit-formation models and is less restrictive - it predicts that one volatility risk is always negatively priced, while the price of the second risk can have both signs. Such a model is consistent with the empirical finding of this paper that one volatility risk is significantly positively priced.

3 Design of the empirical tests

This section describes the construction of the option returns and volatility risk factors, used in testing (2.6).

3.1 Construction of option returns

I construct daily returns on hypothetical options with fixed levels of moneyness and maturity. In particular, the fixed maturity levels allow to focus on possible maturity effects in option returns, which would be blurred if, for example, only options held until expiration are considered. The convenience of working with such constructs is well known and has been exploited in different contexts. From 1993 till 2004 the Chicago Board of Options Exchange was calculating the Volatility Index (VIX) as the implied volatility of a 30-day option, struck at-the money forward. In a research context, Buraschi and Jackwerth (2001) construct 45-day options with fixed moneyness levels close to at-the-money. I follow this approach, and extend it to a number of fixed moneyness levels and maturities, ranging from one month to one year.

Options with predetermined strikes and maturities, most likely, were not actually traded on the exchange on any day in the sample. To find their prices, I apply the following two-step procedure. First, I calibrate an option-pricing model to extract the information contained in the available option prices. Next, this estimated model is used to obtain the prices of the specific options I need. This approach has only recently been made feasible by the advent of models, which are capable of accurately calibrating options in the strike and the maturity dimensions together. Section 3.2 presents three models of this type. Extracting information from available options to price other options is a standard procedure. This is how prices are quoted in over-the-counter option markets - if not currently observed, the option price is derived from other available prices by interpolation or a similar procedure. Also, options that have not been traded on a given day are marked-to-market in traders' books in a similar way[6].

[6] It is possible to avoid the use of an option-pricing model and apply instead some polynomial

Once the model is estimated, any option price can be obtained off it. I consider three levels of moneyness for puts and for calls, and for each one of five maturities, as given in the Table 1. (The table shows the ratios between the actual strikes employed and the at-the-money forward price.)

Table 1 about here

The maturities are one, three, six, nine and twelve months and all strikes are at- or out-of-the-money. To capture the fact that variance increases with maturity, the range of strikes increases accordingly. In this way, for each name I construct 30 time-series of option returns. I estimate volatility risk prices both using all separate time-series and using portfolios, constructed from these series.

Returns to unhedged options are constructed as follows. On each day I calculate option prices on the grid of fixed strikes and maturities. Then I calculate the prices of the same options on the following day - i.e. I keep the strikes but use the following day's parameters and spot price and decrease the time to maturity accordingly. I also take into account the cost of carrying the hedge position to the next day. The daily return on a long zero-cost position in the option is the difference between the second day's option price and the first day's option price with interest:

$$R = O(\theta_2, S_2, K_1, T_2) - O(\theta_1, S_1, K_1, T_1)(1 + r\Delta) \tag{3.1}$$

The indexes 1 and 2 refer to the first and second day and O denotes an option price. S, K, r and T are spot, strike, interest rate and time to maturity respectively, θ is an estimated set of parameters and $\Delta = T_1 - T_2$. Note that as the spot price changes from day to day, the strikes used also change, since the grid of moneyness levels is kept constant. Finally, to make the dollar returns obtained in this way comparable across maturities and names I scale these returns by the option price in the first day.

To calculate returns to delta-hedged options, the delta-hedge ratio is needed as well. I obtain it numerically in the following way: on each day I move up the spot price by a small amount epsilon, calculate the option price at the new spot (everything else kept the same) and divide the difference between the new and old option prices by epsilon. The daily return on a long zero-cost position in the delta-hedged option is the difference between the second day's option price and the first day's option price with interest less delta times the difference between the second day's spot price and the first day's spot

smoothing on the implied volatilities of observed options. While an obvious advantage of this approach is that observed prices are fitted exactly, the downside is that when we need to extrapolate to strikes beyond the range of the observed strikes, this procedure is known to be very inaccurate.

price with interest:

$$R = O(\theta_2, S_2, K_1, T_2) - O(\theta_1, S_1, K_1, T_1)(1 + r\Delta) - \delta(S_2 - S_1(1 + r\Delta)) \tag{3.2}$$

where δ is the delta-hedge ratio and all other parameters are as before.

To make dollar returns comparable across maturities and names, I scale them by the price of the underlying asset. Scaling by the option price is possible, but it disregards the hedging component, which can be much higher than the option one.

Of course, the convenient collection of time-series of option returns comes at the price of daily rebalancing - closing the option position at the previous actual strike and maturity and entering into a new position at a new actual strike (but at the fixed moneyness) and same maturity. By constructing option returns in this way, I assume away the thorny issue of transaction costs. However, I provide an alternative check for the statistical results by considering monthly returns to calendar spreads.

3.2 Data and option-pricing models

All the data I use come from OptionMetrics, a financial research firm specializing in the analysis of option markets. The "Ivy DB" data set from OptionMetrics contains daily closing option prices (bid and ask) for all US listed index and equity options, starting in 1996 and updated quarterly. Besides option prices, it also contains daily time-series of the underlying spot prices, dividend payments and projections, stock splits, historical daily interest rate curves and option volumes. Implied volatilities and sensitivities (delta, gamma, vega and theta) for each option are calculated as well. The comprehensive nature of the database makes it most suitable for empirical work on option markets.

The data sample includes daily option prices of six stock indexes and twenty-two major individual stocks for six full years: 1997 - 2002[7]. Table 2 displays their names and ticker symbols. The 1997-2002 period offers the additional benefit that it can be split roughly in half to obtain a period of steeply rising stock prices (from January 1997 till mid-2000), and a subsequent period of mostly declining stock prices. As a robustness check, results are presented both for the entire period and for the two sub-periods. Table 2 presents also the proportion of three maturity groups in the average daily open interest for at- and out-of-the-money options for each name. It is clear that the longer maturities are well represented, even though the short-maturity group (up to two months) has a somewhat higher proportion in total open interest.

Table 2 about here

[7] 1996 was dropped, since there were much fewer option prices available for this year.

To obtain the option prices needed, I fit a model to the available option prices on each day. The choice among the numerous models that can accurately fit the whole set of options on a given day is of secondary importance in this study. I perform below a limited comparison between three candidate models and pick the one, which is slightly more suitable for my purpose. My main consideration is accuracy of the fit, and I avoid any arguments involving the specifics of the modeled price process. So, both a diffusion-based and a pure-jump model are acceptable, both models with jumps in volatility and in the price process can be used, etc. It turns out that models, which are conceptually quite different, perform equally well for my purpose.

I focus on the following three models: A stochastic volatility with jumps (SVJ) model studied by Bates (1996) and Bakshi et al. (1997), a double-jump (DJ) model, developed in Duffie et al. (2000) and a pure-jump model with stochastic arrival rate of the jumps (VGSA), as in Carr et al. (2003). Appendix C presents some details on the three models. A full-scale comparison between the models is not my purpose here (see Bakshi and Cao (2003) for a recent detailed study). I only compare their pricing accuracy. To do this, I estimate the three models on each day in the sample of S&P500 options (1509 days for the six years). I employ all out-of-the-money options with strike to spot ratio down to 65% for puts and up to 135% for calls, and maturity between one month and one year (140 options per day on average). The main tool for estimation is the characteristic function of the risk-neutral return density, which is available in closed form for all three models (see Appendix B). Following Carr and Madan (1998), I obtain call prices for any parameter set, by inverting the generalized characteristic function of the call price, using the Fast Fourier Transform[8]. I obtain put prices by put-call parity. I then search for the set of parameters, which minimizes the sum of squared differences between model prices and actual prices. The estimation results are as follows: For the DJ model - 47 days with average % error (A.P.E.) above 5% and average A.P.E. of 2.28% in the remaining days. For the SVJ model - 55 days and 2.42% resp. For the VGSA model - 72 days and 2.48% resp. It is reasonable to apply some filter, when working with estimated, not actual option prices. The above estimations show that by discarding days when average A.P.E. is above 5%, fewer than 5% of the data for S&P500 options are lost. I apply this cut-off in all future estimations.

As expected, the richest model (DJ) performs best, but the differences are modest. So, when choosing among the models the advantages of using a more parsimonious model should also be considered. As discussed in Bakshi and Cao (2003), the estimation of large option pricing models on individual names is hindered by data limitations. In their sample, the majority of the 100 most actively traded names on the CBOE have on average less than ten out-of-the-money options per day. That's why they need to

[8] This procedure is strictly valid only for European-style options. Using only at- and out-of-the-money options mitigates the bias introduced from applying this procedure to options on individual stocks which are traded American-style.

pool together options from all days in a week to perform estimations. Such an approach is not feasible in this study, since pooling across the days in a week may hinder the construction of returns to options with fixed moneyness. So I employ only the names with largest number of options per day and further discard days with insufficient number of options. On average, for all names in the sample, fewer than 12 out-of-the-money options are available on 2.4% of the days, fewer than 15 - on 6.8% of the days and fewer than 18 - on 11.3% of the days. Obviously, increasing the degrees of freedom would come at the price of giving up an increasing amount of the data. That is why, mostly a data-related consideration leads me to choose VGSA (which has the fewest parameters), while sacrificing some accuracy. An additional benefit of such a choice is gain in computational speed. I discard all days with fewer than 12 out-of-the-money options.

Table 3 presents the average number of options, used in the estimations for each name, and the percentage errors achieved. I discard days with errors above 5% (usually not more than 4-5% of all days).

Table 3 about here

The errors in the remaining days are around 3% and often less, which is quite satisfactory. This is often within the bid-ask spread, in particular for out-of-the-money options. Armed with the estimated parameters for each day, it is easy to generate model prices and returns at the required strikes and maturities.

3.3 Volatility risk factors

Cross-sectional estimations of risk prices involve regressions of excess returns on measures of respective risks. I construct these risk measures in two steps. First, I calculate, on each day in the sample, proxies for the market's best estimate of market volatility, realized over different subsequent periods (from one month to one year). Second, I calculate the daily changes in these volatility factors to obtain the volatility risk measures (or "volatility risk factors").

The market's estimate of volatility, realized over a given future period is taken to be the price of the volatility swap with the respective maturity. When "the market" is defined to be the S&P500 index and the length of the period is one month, this best estimate is precisely the CBOE's Volatility Index (VIX). VIX is currently calculated via a non-parametric procedure[9] employing all current at- and out-of-the-money short-term options on S&P500. One way to obtain the market's volatility estimates for longer future periods would be to extend this procedure, using options with longer maturities. Alternatively, one can use the estimated model parameters for S&P500 and calculate the

[9] See e.g. Carr and Madan (2001).

standard deviations of the S&P500 risk-neutral distribution at the respective horizons[10]. I apply the second alternative, mostly for computational convenience.

I verify that the two approaches produce very similar results. First, the correlation between VIX and the one-month risk-neutral standard deviation over 1997 - 2002 is 99.1%. Next, I compare the predictive power of the two volatility time-series for realized S&P500 volatility. Following Christensen and Prabhala (1998), I regress realized daily return volatility over 30-day non-overlapping intervals on the two implied volatilities at the beginning of the respective 30-day intervals. Table 4 shows the results for 1997 - 2002 and two sub-periods.

Table 4 about here

The estimates involving the VIX and the risk-neutral standard deviation are almost identical. The similarity is observed both in the entire period and the two sub-periods. This comparison justifies the use of the risk-neutral standard deviation. It also provides an indirect check of the quality of the model-based option prices used in constructing option returns.

I define the volatility risk factors to be the daily changes of the volatility proxies at the respective horizons. I also normalize the volatility risk factors to unit standard deviation, which helps to avoid scaling problems and allows for comparing the prices of different volatility risks.

The calculation of volatility risk factors from option prices is motivated by previous findings that option-implied volatilities at different horizons exhibit quite different behavior, indicating that possibly different risks are embedded in these volatilities: Poterba and Summers (1986) find that the changes in forward short-term implied volatility (which is approximately the difference between short-and long-term volatility) are of the same sign but of consistently smaller absolute value than the changes in current short-term implied volatility. Engle and Mustafa (1992) and Xu and Taylor (1994) find that the volatility of short-term implied volatility is larger and mean-reverts faster than that of long-term implied volatility.

4 Estimation of volatility risk prices

This section presents results on the estimation of market volatility risk prices in the cross-section of expected option returns, as in equation (2.6). It also presents evidence on the economic significance of the difference between the estimated risk prices. For this purpose I employ returns on calendar spreads.

[10] The risk-neutral variance is obtained by evaluating at zero the first and second derivatives of the characteristic function of the risk-neutral distribution at different horizons. The exact form of the second derivative is quite lengthy, but is readily given by any package, implementing symbolic calculations. The volatility proxy is then the square root of this risk-neutral variance.

Table 5 shows summary statistics of the excess option returns used in estimations. Panel A shows average excess returns to unhedged options across the twenty-eight names, for each strike level and each maturity.

Table 5 about here

Observe that there is a wide variation in these average returns to be explained. There is also a clear pattern across maturities - returns to puts invariably increase with maturity, while those to calls decrease with maturity. Overall, average returns to calls are significantly positive, while returns to puts are mostly negative and sometimes not significantly different from zero. Panel B shows average returns to delta-hedged options. The variation in these returns is still considerable. The maturity pattern for puts is preserved and is even more significant than for unhedged options. Interestingly, returns to longer-term calls now tend to be higher than to short-term ones, in contrast to the unhedged case.

4.1 Estimation results

I first estimate the prices of two volatility risk factors with two-step cross-sectional regressions on all individual time-series of excess returns to unhedged options (total of 840 series). I apply the standard procedure of finding the betas on the risk factors at the first step, then regressing, for each day in the sample, excess returns on betas, and finally averaging the second-step regression coefficients and calculating their standard errors.

All regressions involve the market risk, the one-month volatility risk factor and one of the longer term volatility risk factors. In this set-up the one-month risk factor is proxy for short-term volatility risk. This choice is consistent with MacKinlay and Park (2004) who find that monthly volatility exhibits features typical for short-term (transitory) volatility, while three- to six-month volatility represents well long-term (permanent) volatility.

Panel A in Table 6 shows the results of these regressions. Market risk is always positively priced (significant at 5%). One-month volatility risk is always significantly negatively priced,while longer-term volatility risks are always significantly positively priced. These results are to a large extend supported by the cross-sectional regressions involving a single volatility factor. Panel B in Table 6 presents the results for such single-volatility regressions and show that only the one-month volatility risk has a negative price (insignificant), while all other volatility risk prices are positive and mostly significant. Table 6 also shows the importance of a second volatility factor for the explanatory power of the regressions. The numbers in parentheses show, for each combination of risk factors, the proportion of time-series regressions (first pass) with significant alphas. While practically all regressions with one volatility factor have significant alphas (96% in all cases), this proportion dramatically falls to about 15% when a second volatility factor is included. Table 6 also shows the adjusted R^2 in regressing *average* excess returns on betas. For any combination of two volatility factors, R^2 increases by 4-5% compared to the single volatility factor case.

Table 6 about here

These results strongly indicate that, first, two volatility risks are indeed priced in the option market and second, that these risks are of different nature, as evidenced by the different sign of the risk prices. All previous studies of volatility risk pricing in the option market use relatively short-term options (maturity about one month) and mostly find a negative risk price. So, my finding of a negatively priced short-term volatility risk is consistent with previous empirical studies. However, a positively priced long-term volatility risk has not been identified before. Table 6 also demonstrates the ability of the three-factor model to capture the variance in option excess returns.

One possible concern with the results in Table 6 relates to the correlations between the market and the volatility risk factors. This correlation is well known to be negative and typically high in magnitude as discussed in Section 2.1. For the five volatility risk fdactors used in the estimations the correlation is -0.60% or less. Besides, the volatility risk factors are highly positively correlated among themselves (50-90% in the sample). To eliminate the possible effect of these correlations, I run the regressions also with orthogonal volatility risk factors - I use the component of the one-month volatility risk which is orthogonal to market returns and the component of the longer-term volatility risk which is orthogonal to both the market and the one-month orthogonal volatility risk. Table 7 compares the results involving the original (raw) volatility risks with these "orthogonal" volatility risks.

Table 7 about here

When orthogonal volatility risk factors are used (Panel B), the results are qualitatively the same. The significance of the negative price of the one-month factor is sometimes lower. Note that the price of risk in the orthogonal long-term factor is much higher in magnitude compared to the raw factor. The estimations with volatility risk factors, orthogonal to market returns also indicate that specific volatility risks are priced in the option market. Investors recognize volatility risks beyond those, due to the negative correlation between changes in volatility and market returns.

The above test can be subject to several concerns. First, a more precise evaluation of the explanatory power of the model can be performed using the joint distribution of the errors. However, given the large amount of time series involved, this task is computationally quite demanding, as it involves the computation and inversion of the covariance matrix of the residuals in the time-series regressions. Second, the standard errors of the estimated parameters are not corrected for heteroskedasticity. This may be an issue when dealing with option data and volatility risks, given the persistence in volatility. Third, unhedged options mix the exposure to the risk in the underlying asset and to volatility risk. It can be argued that investors, seeking exposure to volatility risk will hedge away the risk in the underlyer. So, the price of volatility risk may be better reflected in returns to hedged option. Fourth, the different names are given equal weight, even though their relative importance is highly unequal - for example options on the S&P500 amount to almost half the value of all options in the sample.

To address these concerns I apply a second test where I first, consider delta-hedged instead of unhedged options, second, apply GMM for the estimation and third, reduce the number of asset-return series by forming option portfolios. Delta-hedging allows for more precise exposure to volatility risk. GMM handles the heteroskedasticity of errors and allows to test for all errors being jointly equal to zero. The portfolios allow for an efficient implementation of GMM and account to an extent for the relative importance of different options. Forming portfolios addresses one deficiency of the returns data as well. Because of insufficient out-of-the-money options on certain days and because sometimes the error of estimation has been too high, there are missing observations for certain days for each returns time-series. Since the omissions are relatively few and they come at different days for different names, having several names in a portfolio leaves no missing data in the aggregated returns series.

To form the portfolios I sort the names in the sample according to their average implied volatility (see Table 2). Each portfolio belongs to one of the five maturity groups and one of five volatility quintiles (a total of twenty-five portfolios). The different strikes for each name are weighted by the average open interest for the closest available strikes in the data. For example, I find the proportion of the closest to at-the-money puts on each date and assign the average of these proportions across all days to be the weight of at-the-money puts. I proceed in the same way with the other moneyness levels, both for puts and for calls. The different names within a portfolio are weighted by the average option value for the name, where the average is taken again across all days in the six-year period.

Table 8 shows the average excess returns on the portfolios (a total of twenty five) over 1997 - 2002. As in Table 3 the numbers are in percent and on a monthly basis. The columns show portfolios arranged from the lowest to the highest average implied volatility of the components.

Table 8 about here

Compared to Table 5, the weighted returns (i.e. the portfolio returns) tend to be much lower. There are still portfolios with positive returns, but fewer and with smaller absolute returns. Obviously the larger and less volatile names (in particular the indexes) tend to have lower returns. What is preserved however is the maturity effect - returns to longer-maturity options tend to be higher.

Tables 9 and 10 contain the main result of this paper. Table 9 shows volatility risk prices from estimations with two volatility risk factor. GMM estimations with ten Newey-West lags are reported[11]. Results are reported both for the entire six-year period and separately for 1997 to mid-2000 and from mid-2000 to 2002.

Table 9 about here

The one-month volatility factor is always included; the longer-term factors are included both in their raw form, and only with their component orthogonal to the one-month volatility risk factor. Using the orthogonal component does not change the remaining estimates. In all cases the market risk is not priced. This can be expected given delta-hedging. The price of one-month volatility risk is negative and highly significant, except for the second sub-period. However, all other volatility risk prices are positive. The prices of the raw factors are marginally significant (z-statistics about 1.5 - 1.6). Note that the significance is higher for the longer maturity raw factors. However, the risk prices of the orthogonal components are all significant (with two exceptions in the second sub-period). The p-values in all cases are very high (typically 80% or more). The adjusted R^2 in regressing average returns on betas is typically high (0.80 or more), except for the second sub-period. We have thus strong indication that two volatility factors explain most of the variation in expected option returns in the data sample.

Table 10 shows volatility risk prices from estimations with one volatility risk factors, which markedly contrast with the two-factor case.

Table 10 about here

The price of market risk is again insignificant. The one-month and three-month volatility risks for the whole period are significantly negatively priced. The prices of longer-term volatility risks are all negative, but not significant. Note that both the significance levels and the absolute magnitude of the volatility risk prices steadily decrease as maturity increases. The same pattern is exactly repeated in the first sub-period. The second sub-period presents mostly insignificant estimates. The table also shows p-values for the chi-squared test for all pricing errors being jointly zero. The entire period and the first sub-period have high p-values for the estimation with the one-month factor (24% and 30% resp.). For longer maturities the p-values decrease sharply. The second sub-period is again different, showing very high p-values (above 60%) for all maturities. The relation between average returns and betas is now weaker (adjusted R^2 about 0.60 or less, and even negligible in the second sub-period).

The results in Tables 9 and 10 clearly show that long-term volatilities contain two

[11] Five and twenty lags were also used, producing very similar results which are not reported.

separate risk components with different prices. Including both short- and long-term volatility in the regressions helps disentangle these two risk components. The second risk, additional to short-term volatility risk is positively priced. The magnitude of the price of this second risk is typically equal or higher than that of the short-term risk price[12].

What is the actual maturity of the short-term volatility risk? So far, the one-month volatility factor is assumed to be the short-term one. However, previous studies have considered various frequencies, sometimes much shorter than a month. It is then possible that the true short-term factor is of much lower maturity, and the one-month factor also mixes two risks. I address this issue by assuming that the absolute value of market returns is proxy for very short-term (one-day) volatility. I run regressions with absolute market returns and each of the former volatility risk factors. If the true maturity of short-term risk is well below one month, this will be reflected in significant estimates of the price of one-day volatility risk. In this case the one-day volatility risk can also be expected to help disentangle the risks in the longer term volatilities. The high explanatory power of two volatility factors for option returns (Table 9) should also be preserved. On the other hand, if the true maturity of short-term risk is well above one day, the regressions results should resemble those of the one-volatility factor case (Table 10).

Table 11 presents the results of regressions involving absolute market returns.

Table 11 about here

These regressions do not support the hypothesis of a very short maturity of the short term-risk. The estimated risk prices for one- to twelve-month volatilities and their significance levels are very close to the one-factor case (Table 10). The explanatory power of the regressions is also similar to that in the one-factor case. The absolute value of market returns indeed has a negative and often significant price, but given the other estimates this is not likely to reflect a separate very-short term source of volatility risk.

In summary, I find that two market volatility risks are significantly priced in the sample of excess option returns. The model with two volatility risk factors has little pricing error. While the short-term factor, embedded in close to one-month implied volatility is negatively priced, the long-term factor which can be extracted from longer term-volatilities is positively priced. Next I consider returns to calendar spreads to investigate whether the difference between the prices of short- and long-term volatility risks is also economically significant.

[12] The tests reported in Tables 8 and 9 were also run with the components of short- and long-term volatilities, orthogonal to market returns. The results were very close, which can be expected, given the insignificant estimates for market risk price.

4.2 Evidence from calendar spreads

A long calendar spread is a combination of a short position in an option with short maturity and a long position in an option on the same name, of the same type and with the same strike, but of a longer maturity. These spreads are similar to options, in the sense that the possible loss is limited to the amount of the initial net outlay. The results from cross-sectional regressions on unhedged options reported in Table 6 showed a positively priced market risk and long-term volatility risk and a negatively priced short-term risk. Expected returns to calendar spreads then have two components - one reflecting the market risks in the two options in the spread and another related to the two volatility risks. The sign and magnitude of the first component should differ across types of options (calls or puts) and moneyness (see e.g. Coval and Shumway (2001)). Results in the previous section imply that the sign of the second component should be unambiguously positive (the position is short the short-term risk and long the long-term one). I do not derive here a formal relation between the two components, but verify that for all moneyness ranges and for both option types (put and calls) the expected returns to calendar spreads are positive. This demonstrates that the component related to volatility risk is always positive and dominates the market risk component.

Since the results so far only concern market volatility risks, I should strictly focus here only on calendar spreads written on the market. However, given the high explanatory power of market volatility risk factors for option returns (Table 9), it can also be expected that calendar spreads written on individual names reflect the difference in the pricing of market volatility risks. So, I consider spreads written on all names in the sample as well.

I use all options in the data set, which allow to calculate the gain of a position in calendar spread. For each name I record, at the beginning of each month all couples of options of the same type and strike and with different maturities, for which prices are available at the end of the month as well. In each spread I use a short-term option of the first available maturity above 50 days. Possible liquidity problems when reversing the position at the end of the month are thus avoided. In this way I replicate a strategy, which trades only twice every month - on opening and closing the spread position. While transaction costs are still present, such trades are definitely feasible. I calculate returns on spreads where the long term option is of the second or third available maturities above 50 days. The results are very similar and I only report results for the second maturity.

Table 12 shows average returns to calendar spreads written on the market (S&P500) and on all names in the sample. Separately are shown average returns for different ranges of moneyness for puts and calls, and for different periods.

Table 12 about here

On average, spreads on puts gain an impressive 20% monthly, while those on calls - about 12%. Average returns on individual names are slightly lower than those on the market alone. Spread returns in different moneyness ranges vary considerably, but are all positive. Table 12 also show that the Sharpe ratios of calendar spreads are typically about 30-40%, going as high as 100% in one case. Transaction costs would reduce these numbers, but still the differential pricing of volatility risks is very pronounced. Spread returns thus show that the different prices of short- and long-term volatility risks are not only statistically significant, but economically significant as well.

5 Conclusion

A number of volatility-related financial products have been introduced in recent years. Derivatives on realized variance and volatility have been actively traded over the counter. In 2004 the CBOE Futures Exchange introduced futures on the VIX and on the realized three-month variance of the S&P500 index. The practitioners' interest in volatility products has been paralleled by academic research of volatility risk, mostly focused on the risk embedded in a single volatility factor.

This paper complements previous studies of volatility risk by presenting evidence that two implied-volatility risks are priced in a cross-section of expected option returns. I find that the risk in short-term volatility is significantly negatively priced, while another source of risk, orthogonal to the short-term one and embedded in longer-term volatility is significantly positively priced. I show further that the difference in the pricing of short- and long-term volatility risks is also economically significant: I examine returns on long calendar spreads and find that, on average, spreads gain up to 20% monthly.

The estimations for the two sub-periods reveal considerable differences in the parameters, indicating that an extension to time-varying betas and risk prices is justified. The robustness of the findings in this paper to the choice of an option data set and an option-pricing model in constructing option returns can also be examined.

The differential pricing of volatility risks has implication for the modeling of investors' utility. Previous research has found evidence for utility functions over wealth which have both concave and convex sections (Jackwerth (2000), Carr et al. (2002)). It is interesting to explore their results by employing utility functions with more than one arguments and possibly different volatility risks.

Another implication of the findings in this paper relates to the use of options in risk management. It has been argued that firms sometimes face risks which are bundled together in a single asset or liability (e.g. Schrand and Unal (1998)). In this case they can use derivatives to allocate their total risk exposure among multiple sources of risk. This paper suggests that derivatives themselves reflect multiple risks. How do firms chose among derivatives incorporating multiple risks is still to be studied.

References

[1] Abel, A. (1990), "Asset prices under habit formation and catching up with the Joneses", *American Economic Review* 80, 38-42.

[2] Alizadeh, S., M. Brandt and F. Diebold (2002), "Range-based estimation of stochastic volatility models", *Journal of Finance* 57, 1047-1091.

[3] Andersen,T. and T. Bollerslev (1997), "Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns", *Journal of Finance* 52, 975-1005.

[4] Ang, A., R. Hodrick, Y. Xing and X. Zhang (2004),"The cross-section of volatility and expected returns", Working paper, Columbia University.

[5] Bansal, R. and A. Yaron (2004), "Risks for the long run", *Journal of Finance* 59, 1481-1509.

[6] Bakshi, G. and C. Cao (2003), "Risk-neutral kurtosis, jumps and option pricing: Evidence from 100 most actively traded firms on the CBOE", Working paper, University of Maryland.

[7] Bakshi, G. and N. Kapadia (2003), "Delta-hedged gains and the negative volatility risk premium", *Review of Financial Studies* 16, 101-143.

[8] Bekaert, G. and G. Wu (2000), "Asymmetric volatilities and risk in equity markets", *Review of Financial Studies* 13, 1-42.

[9] Benzoni, L. (2001) "Pricing options under stochastic volatility", Working paper, University of Minnesota.

[10] Black, F. (1976), "Studies of stock price volatility changes", *Proceedings of the 1976's meeting of the Business and Economic Statistics Section*, American Statistical Association, 177-181.

[11] Buraschi, A. and J. Jackwerth (2001), "The price of a smile: hedging and spanning in option markets", *Review of Financial Studies* 14, 495-527.

[12] Carr, P. and D. Madan (1999), "Option pricing and the Fast Fourier transform", *Journal of Computational Finance* 4, 61-73.

[13] Carr, P. and D. Madan (2001), "Towards a theory of volatility trading", in *Volatility* R. Jarrow (ed.) Risk publications.

[14] Carr, P., H. Geman, D. Madan and M. Yor (2002), "The fine structure of asset returns", *Journal of Business* 75, 305-332.

[15] Carr, P., H. Geman, D. Madan and M. Yor (2003), "Stochastic volatility for Levy processes", *Mathematical Finance* 13, 345-382.

[16] Carr, P. and L. Wu (2004), "Time-changed Levy processes and option pricing", *Journal of Financial Economics* 71, 113-141.

[17] Carr, P. and L. Wu (2004), "Variance risk premia", Working paper, Baruch College.

[18] Chernov, M. and E. Ghysels (2000), "A study towards a unified approach to the joint estimation of objective and risk-neutral measures for the purpose of option valuation", *Journal of Financial Economics* 56, 407-458.

[19] Chernov, M., A. Gallant, E. Ghysels and G. Tauchen (2003), "Alternative models for stock price dynamics", *Journal of Econometrics* 116, 225-257.

[20] Cochrane, J. (2001), *Asset Pricing*, Princeton University Press.

[21] Coval, J. and T. Shumway (2001), "Expected option returns," *Journal of Finance* 56, 983-1009.

[22] Christensen, B. and N. Prabhala (1998), "The relation between implied and realized volatility", *Journal of Financial Economics* 50, 125-150.

[23] Doran, J. and E. Ronn (2003), "On the market price of volatility risk", Working paper, University of Texas.

[24] Duffie, D., J. Pan and K. Singleton (2000) "Transform analysis and asset pricing for affine jump-diffusions", *Econometrica* 68, 1343-1376.

[25] Engle, R. and G. Lee (1998), "A long-run and short-run component model of stock return volatility", in R. Engle and H. White (eds.), *Co-integration, causality and forecasting: A Festschrift in Honor of Clive W. J. Granger*, Oxford University Press, 475-497.

[26] Engle, R. and C. Mustafa (1992), "Implied ARCH models from option prices", *Journal of Econometrics* 52, 289-311.

[27] Gali, J. (1994), "Keeping up with the Joneses: consumption externalities, portfolio choice and asset prices", *Journal of Money, Credit and Banking* 26, 1-8.

[28] Gallant, A., C. Hsu, and G. Tauchen (1999), "Using daily range data to calibrate volatility diffusions and extract the forward integrated variance", Working paper, Duke University.

[29] Heston, S. (1993), "A closed-form solution for options with stochastic volatility and application to bond and currency options", *Review of Financial Studies* 6, 327-343.

[30] Jones, C. (2003), "The dynamics of stochastic volatility", *Journal of Econometrics* 116, 181-224.

[31] Jackwerth, J. (2000), "Recovering risk aversion from option prices and realized returns", *Review of Financial Studies* 13, 433-451.

[32] Liesenfeld, R. (2001) "A generalized bivariate mixture model for stock price volatility and trading volume", *Journal of Econometrics* 104, 141-178.

[33] MacKinlay, C. and J. Park (2004), "The relationship between expected risk premium and conditional permanent and transitory volatility", Working paper, Wharton School.

[34] Madan, D., P. Carr and E. Chang (1998), "The Variance Gamma process and option pricing", *European Finance Review* 1, 79-105.

[35] Muller, U., M. Dacorogna, R. Dave, R. Olsen and O. Pictet (1997), "Volatilities of different time resolutions - analyzing the dynamics of market components", *Journal of Empirical Finance* 4, 213-239.

[36] Pan, J. (2002), "The jump-risk premia implicit in options: Evidence from an integrated time-series study", *Journal of Financial Economics* 63, 3-50.

[37] Poterba, J. and L. Summers (1986), "The persistence of volatility and stock-market fluctuations", *American Economic Review* 76, 1142-1151.

[38] Schrand, C and H. Unal (1998), "Hedging and coordinated risk management", *Journal of Finance* 53, 979-1013.

[39] Tauchen, G. (2004), "Stochastic volatility in general equilibrium", Working paper, Duke University.

[40] Xu, X. and S. Taylor (1994),. "The term structure of volatility implied by foreign exchange options", *Journal of Financial and Quantitative Analysis* 29, 57-74.

A Proof of Proposition 1

Expectations are taken under the statistical measure. The pricing kernel process Λ_t is obtained by conditioning on the realizations of the volatility σ^S. Let $\overline{E_t}[\cdot]$ denote expectation, conditional on the volatility path.

$$C_T = C_t \exp\left(\int_t^T D_s^S ds + \int_t^T \sqrt{\sigma_s^S} dW_s^C - \frac{1}{2}\int_t^T \sigma_s^S ds\right) \tag{A.1}$$

$$\begin{aligned}
\Lambda_t &= E_t\left[C_t^{-\lambda} \exp\left(-\lambda \int_t^T D_s^S ds - \lambda \int_t^T \sqrt{\sigma_s^S} dW_s^S + \frac{1}{2}\lambda \int_t^T \sigma_s^S ds\right)\right] \\
&= C_t^{-\lambda} \exp\left(-\lambda \int_t^T D_s^S ds\right) E_t\left[\exp\left(\frac{1}{2}\lambda \int_t^T \sigma_s^S ds\right) \overline{E_t}\left[\exp\left(-\lambda \int_t^T \sqrt{\sigma_s^S} dW_s^S\right)\right]\right] \\
&= C_t^{-\lambda} \exp\left(-\lambda \int_t^T D_s^S ds\right) E_t\left[\exp\left(\frac{\lambda(\lambda+1)}{2} \int_t^T \sigma_s^S ds\right)\right] \qquad (A.2)
\end{aligned}$$

To evaluate the expectation in (A.2) we use the fact that the characteristic function of the integral of a CIR variable is known in closed form. Let $\tau = T - t$ and $Y_\tau = \int_t^T \sigma_s^S ds$. Then the characteristic function of Y_τ is:

$$\phi_{Y_\tau}(u) = E_t\left[e^{iuY_\tau}\right] = A(\tau, u) \exp\left[B(\tau, u)\sigma_t^C\right] \tag{A.3}$$

$$\begin{aligned}
A(\tau, u) &= \frac{\exp(\frac{\kappa^2\theta\tau}{\eta^2})}{\left(\cosh(\frac{\gamma\tau}{2}) + \frac{\kappa}{\gamma}\sinh(\frac{\gamma\tau}{2})\right)^{\frac{2\kappa\theta}{\eta^2}}} \\
B(\tau, u) &= \frac{2iu}{\kappa + \gamma\coth(\frac{\gamma\tau}{2})} \qquad \text{and} \qquad \gamma = \sqrt{\kappa^2 - 2\eta^2 iu}
\end{aligned}$$

On evaluating this characteristic function at $u = -i\frac{\lambda(\lambda+1)}{2}$ it follows from (A.2) that the stochastic discount factor is:

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda\frac{dC_t}{C_t} - \frac{\lambda(\lambda+1)}{\kappa + \gamma\coth(\frac{\gamma\tau}{2})} d\sigma_t^C \tag{A.4}$$

whereby terms of order dt in the differentiation of Λ_t are ignored.

B Volatility risk prices in a two-factor model

This appendix presents a stylized model which is consistent with the empirical findings of the paper. The model is an extension of the basic model in (1)-(3), which includes a second stochastic volatility. The stochastic discount factor derived for the extended model corresponds closely to the two-factor model (2.6) tested empirically in this paper. For the extended model I show that one volatility risk has always a negative price, while the price of the second volatility risk can be of both signs. In particular, this risk price is positive (as found empirically above) when the utility function is convex in the second volatility factor.

In the spirit of habit-formation models, I assume an economy where utility is of the form:

$$U_T = \frac{\left(\frac{C_T}{H_T}\right)^{1-\lambda}}{1-\lambda} \tag{B.1}$$

and the pricing kernel process is:

$$\Lambda_t = E_t\left[\left(\frac{C_T}{H_T}\right)^{-\lambda}\right] \tag{B.2}$$

where C denotes an aggregate consumption good, and λ is the risk aversion coefficient. H is interpreted as a time-varying habit in a multiplicative form, as introduced by Abel (1990) and Gali (1994). In models of this type utility does not depend on the absolute level of consumption, but on the level of consumption relative to a benchmark (habit). The habit is usually related to past consumption. Here I also allow for randomness in habit. Assume the following dynamics for H and C and their respective volatility processes:

$$\frac{dH_t}{H_t} = D_t^H dt + \sqrt{\sigma_t^H} dW_t^H \tag{B.3}$$

$$\frac{dC_t}{C_t} = D_t^C dt + \sqrt{\sigma_t^C} dW_t^C + \beta\sqrt{\sigma_t^H} dW_t^H \tag{B.4}$$

$$d\sigma_t^C = k^C(\theta^C - \sigma_t^C)dt + \eta^C\sqrt{\sigma_t^C} dW_t^1 \tag{B.5}$$

$$d\sigma_t^H = k^H(\theta^H - \sigma_t^H)dt + \eta^H\sqrt{\sigma_t^H} dW_t^2 \tag{B.6}$$

All four Brownian motions W_t^H, W_t^C, W_t^1 and W_t^2 are assumed to be independent. The drift of habit growth (D_t^H) can be a function of past consumption as in the "catching up with the Joneses" versions of habit models; it is not modeled explicitly as before. What is essential is the separate source of randomness in habit - W_t^H, with stochastic volatility $\sqrt{\sigma_t^H}$ ("habit volatility"). Such a source of randomness reflects the notion that habit also depends on some current variables, similar to the "keeping up with the Joneses" versions of habit-formation models. The drift D_t^C is not modeled explicitly as before. A key feature of the the model is that it allows for random changes in habit to affect the dynamics of the consumption good (β is a sensitivity parameter, so it should take values between zero and one). We now prove the following:

Proposition 2 *The stochastic discount factor ζ_t in the economy with two volatility factors described in (B.1)-(B.6) is given by*

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda\frac{dC_t}{C_t} - \lambda\frac{dH_t}{H_t} - B^C(\tau, u^C)d\sigma_t^C - B^H(\tau, u^H)d\sigma_t^H \tag{B.7}$$

where

$$B^C > 0 \quad \textit{and}$$
$$B^H > 0 \quad \textit{if} \quad \frac{\lambda - 1}{\lambda + 1} < \beta < 1 \tag{B.8}$$

Proof. Expectations are taken under the statistical measure. For brevity in notation let:

$$\overline{C_\tau} = \exp\left(\int_t^T D_s^C ds\right) \quad \text{and} \quad \overline{H_\tau} = \exp\left(\int_t^T D_s^H ds\right)$$

Condition on the realizations of σ^H and σ^C, and obtain for H_T and C_T the following:

$$H_T = H_t\overline{H_\tau}\exp\left(\int_t^T \sqrt{\sigma_s^H}dW_s^H - \frac{1}{2}\int_t^T \sigma_s^H ds\right) \tag{B.9}$$

$$C_T = C_t\overline{C_\tau}\exp\left(\int_t^T \sqrt{\sigma_s^C}dW_s^C + \beta\int_t^T \sqrt{\sigma_s^H}dW_s^H\right)\exp\left(-\frac{1}{2}\int_t^T \sigma_s^C ds - \frac{1}{2}\beta^2\int_t^T \sigma_s^H ds\right) \tag{B.10}$$

The pricing kernel is:

$$\Lambda_t = C_t^{-\lambda}\overline{C_\tau}^{-\lambda}H_t^\lambda\overline{H_\tau}^\lambda E_t\left[L\overline{E_t}\left[M\right]\right] \tag{B.11}$$

where

$$L = \exp\left(\frac{1}{2}\lambda\left(\int_t^T \sigma_s^C ds + (\beta^2 - 1)\int_t^T \sigma_s^H ds\right)\right)$$

$$M = \exp\left(-\lambda\left(\int_t^T \sqrt{\sigma_s^C} dW_s^C + (\beta-1)\int_t^T \sqrt{\sigma_s^H} dW_s^H\right)\right)$$

Take first the conditional expectation in (29):

$$\begin{aligned}
\Lambda_t &= C_t^{-\lambda}\overline{C_\tau}^{-\lambda} H_t^{\lambda}\overline{H_\tau}^{\lambda} E_t\left[L\exp\left(\frac{1}{2}\lambda^2\left(\int_t^T \sigma_s^C ds + (\beta-1)^2\int_t^T \sigma_s^H ds\right)\right)\right] \\
&= C_t^{-\lambda}\overline{C_\tau}^{-\lambda} H_t^{\lambda}\overline{H_\tau}^{\lambda} E_t[N] \qquad \text{(B.12)}
\end{aligned}$$

where

$$N = \exp\left(\frac{1}{2}\lambda(\lambda+1)\int_t^T \sigma_s^C ds - \frac{1}{2}\lambda(1-\beta)\left[1+\beta-\lambda(1-\beta)\right]\int_t^T \sigma_s^H ds\right)$$

Using the independence of the two volatility processes and (A.3), the pricing kernel process is:

$$\Lambda_t = C_t^{-\lambda}\overline{C_\tau}^{-\lambda} H_t^{\lambda}\overline{H_\tau}^{\lambda} A^C(\tau,u^C)\exp(B^C(\tau,u^C)\sigma_t^C)A^H(\tau,u^H)\exp(B^H(\tau,u^H)\sigma_t^H) \qquad \text{(B.13)}$$

where superscripts C and H denote parameters related to consumption and habit respectively, and

$$\begin{aligned}
u^C &= -i\frac{\lambda(\lambda+1)}{2} \\
u^H &= i\frac{1}{2}\lambda(1-\beta)\left[1+\beta-\lambda(1-\beta)\right] \\
B^C(\tau,u^C) &= \frac{\lambda(\lambda+1)}{\kappa^C+\gamma^C\coth(\frac{\gamma^C\tau}{2})} > 0 \qquad \text{(B.14)} \\
B^H(\tau,u^H) &= -\frac{\lambda(1-\beta)\left[1+\beta-\lambda(1-\beta)\right]}{\kappa^H+\gamma^H\coth(\frac{\gamma^H\tau}{2})}
\end{aligned}$$

The stochastic discount factor is then:

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda\frac{dC_t}{C_t} - B^C(\tau,u^C)d\sigma_t^C - \lambda\frac{dH_t}{H_t} - B^H(\tau,u^H)d\sigma_t^H \qquad \text{(B.15)}$$

and $B^H(\tau,u^H) > 0$ if $\frac{\lambda-1}{\lambda+1} < \beta < 1$ ■

The first two terms are exactly as in the case with one volatility factor and without habit. The third term reflects a negative habit risk price. Given that an increase in habit decreases utility, any asset positively correlated with the change in habit has high payoffs in low-utility states, which provides the intuition for the negative price of habit risk. The last term reflects the price of "habit volatility" risk. Assume further that consumption equals dividends and that the market price-to-dividend ratio is constant. So, the market risk price is positive and the prices of the two market volatility risks have the signs of the risk prices of σ_t^C and σ_t^H respectively. Then the stochastic discount factor can also be written as:

$$\zeta_t = -\frac{d\Lambda_t}{\Lambda_t} = \lambda\frac{dS_t}{S_t} - \lambda\frac{dH_t}{H_t} + \lambda^{\sigma} d\sigma_t + \lambda^{\sigma^H} d\sigma_t^H \tag{B.16}$$

$\lambda^{\sigma} < 0$ as follows from (B.14) and (B.15) and the sign of λ^{σ^H} is determined from (B.8).

What is the intuition for condition (B.8)? We now show that this condition is closely related to the concavity or convexity of the utility function (B.1) in habit. Denote for brevity

$$\widetilde{C_T} = C_t\overline{C_\tau}\exp\left(\int_t^T \sqrt{\sigma_s^C} dW_s^C - \frac{1}{2}\int_t^T \sigma_s^C ds\right)$$

Re-write (B.10) as:

$$\begin{aligned} C_T &= \widetilde{C}_T \exp\left(\beta\int_t^T \sqrt{\sigma_s^H} dW_s^H - \frac{1}{2}\beta^2\int_t^T \sigma_s^H ds \pm \frac{1}{2}\beta\int_t^T \sigma_s^H ds\right) \\ &= \widetilde{C}_T H_t^{-\beta}\overline{H_\tau}^{-\beta}\exp\left(\frac{1}{2}\beta(1-\beta)\int_t^T \sigma_s^H ds\right) H_T^{\beta} \end{aligned} \tag{B.17}$$

Then utility at time T is:

$$U_T = \frac{1}{1-\lambda}\left[\widetilde{C}_T H_t^{-\beta}\overline{H_\tau}^{-\beta}\exp\left(\frac{1}{2}\beta(1-\beta)\int_t^T \sigma_s^H ds\right)\right]^{1-\lambda} H_T^{(\beta-1)(1-\lambda)} \tag{B.18}$$

This function is convex in H_T when $\frac{\lambda-2}{\lambda-1} < \beta < 1$. While this condition is not equivalent to (B.8), it can be easily checked that for $\lambda > 3$ the positive price of habit volatility risk follows from the convexity in habit of the utility function. Empirical studies show that $\lambda > 3$ is a plausible range for the values of the risk aversion parameter. So, intuition for the the price of "habit volatility" risk can be provided, as in the single volatility risk case by the shape of the utility function - concavity in habit reflects a decrease in expected utility when "habit volatility" increases, hence the desirability of assets correlated with changes in "habit volatility" and the negative price of "habit volatility" risk. Exactly the opposite argument applies when the utility function is convex in habit, so the price of "habit volatility" risk is positive in this case.

The model presented above is illustrative in nature. It leaves unspecified some important components, in particular the form of the two drifts D_t^H and D_t^C. It also assumes a constant price-dividend ratio. Still, it points to the essential role that concavity / convexity of utility functions can have in modeling volatility risks.

Utility functions which exhibit both concavity and convexity have been found empirically in different contexts. From results in Jackwerth (2000) it follows that investors' utility from wealth (proxied by a market index) has been changing over time. In particular after 1987 it has exhibited a convex shape over certain ranges of market moves. In a similar spirit, Carr et al. (2002) estimate marginal utility over instantaneous market moves (jumps) of different size and show that it has both decreasing and increasing sections over different ranges of jumps. The increasing sections correspond to convex utility. It is interesting to explore their results in the context of utility functions with two arguments and possibly two volatility risks.

Now compare (B.16) with (2.6), estimated above. Note that (2.6) lacks a term corresponding to $\frac{dH_t}{H_t}$. However, this lack is reflected only in the pricing errors, but not in the prices of volatility risks (for uncorrelated risk factors). Besides, with excess returns and normalized risk factors the risk prices in the two models are equivalent, as discussed above. It follows that the extended model is consistent with the empirical findings of this paper

C Option pricing models

This appendix presents details on the three option pricing models compared in Section 3:

1. Stochastic volatility and jumps model

Bates (1996) and Bakshi et al. (1997) consider an eight-parameter model for the log price with stochastic volatility and jumps (SVJ):

$$\begin{aligned} dS_t &= (r-\lambda)S_t dt + \sqrt{V_t} S_t dW_t^1 + J_y S_t dq_t^y \\ dV_t &= (\theta - kV_t)dt + \sigma\sqrt{V_t} dW_t^2 \end{aligned} \qquad \text{(C.1)}$$

W_t^1 and W_t^2 are standard Brownian motions with correlation ρ,
J_y is log-normal with mean $\mu_y$and variance σ_y^2

q_t^y is a Poisson process with arrival rate λ_y
r is interest rate and λ is the jump-compensator

Let: $a = u^2 + iu + 2iu\mu_y\lambda_y - 2\lambda_y(e^{iu\mu_y - \frac{u^2\sigma_y^2}{2}} - 1)$
$b = iu\sigma\rho - k$
$\gamma = \sqrt{b^2 + a\sigma^2}$
$c = \theta\left[(\gamma + b)t + 2\log(1 - \frac{\gamma+b}{2\gamma}(1 - e^{-\gamma t}))\right]$
$d = 2\gamma - (\gamma + b)(1 - e^{-\gamma t})$
$B = -\frac{a(1-e^{-\gamma t})}{d}$

The characteristic function of the log-price at horizon t under the SVJ process is:

$$E^Q\left[e^{iu\log(S_t)}\right] = e^{iu\log(S_0)+iurt-\frac{c}{\sigma^2}+BV_0} \tag{C.2}$$

where V_0 is instantaneous variance. The expectation is taken here under the risk-neutral measure Q.

2. Double jumps model

Duffie et al. (2000) develop a flexible model specification, which also adds jumps in volatility. I use here one of their double-jump (DJ) models:

$$\begin{aligned} dS_t &= (r - \lambda)S_t dt + \sqrt{V_t}S_t dW_t^1 + J_y S_t dq_t^y \\ dV_t &= (\theta - kV_t)dt + \sigma\sqrt{V_t}dW_t^2 + J_v dq_t^v \end{aligned} \tag{C.3}$$

All variables, except for J_v and q_t^v have the same meaning as above. q_t^v is a second Poisson process with arrival rate λ_v. The sizes of the jumps are exponentially distributed with mean μ_v

Let: $a = u^2 + iu$
$b = iu\sigma\rho - k$
$\gamma = \sqrt{b^2 + a\sigma^2}$
$c = \theta\left[(\gamma + b)t + 2\log(1 - \frac{\gamma+b}{2\gamma}(1 - e^{-\gamma t}))\right]$
$d = 2\gamma - (\gamma + b)(1 - e^{-\gamma t})$
$f = \frac{\lambda_y e^{\mu_y + \frac{\sigma_y^2}{2}} + \lambda_v}{\lambda_y + \lambda_v} - 1$
$g = te^{iu\mu_y - \frac{u^2\sigma_y^2}{2}}$
$h = \frac{\gamma - b}{\gamma - b + a}t - \frac{2\mu_v a}{\gamma^2 - (b - \mu_v a)^2}\log(1 - \frac{\gamma + b - \mu_v a}{2\gamma}(1 - e^{-\gamma t}))$

$A = -(\lambda_y + \lambda_v)(1 + iuf)t + \lambda_y g + \lambda_v h$
$B = -\frac{a(1-e^{-\gamma t})}{d}$

The characteristic function of the log-price at horizon t under the DJ process is:

$$E^Q\left[e^{iu\log(S_t)}\right] = e^{iu\log(S_0)+iurt-\frac{c}{\sigma^2}+A+BV_0} \tag{C.4}$$

3. VGSA model

Carr and Wu (2003) provide a general study of the financial applications of time-changed Levy processes. They show that most of the stochastic processes employed as models of asset returns, including SVJ and DJ, belong to this class. The third model I consider here is also based on a process of this class. The VGSA process (Carr et al. (2003)) is a six-parameter pure-jump process with stochastic arrival rate of jumps. (Stochastic arrival is an analogue of stochastic volatility for pure-jump processes.) VGSA is introduced in two steps, each step being an explicit time-change of a Levy process.

At the first step a Brownian motion with drift, denoted as

$$b(t;\theta,\sigma) = \theta t + \sigma W_t \tag{C.5}$$

where W_t is standard Brownian motion, is evaluated at a time given by an independent gamma process $\gamma(t;1,\nu)$ with unit mean rate and variance rate ν. The process, obtained in this way is the Variance Gamma (VG) process (Carr and Madan (1998)):

$$VG(t;\sigma,\nu,\theta) = b(\gamma(t;1,\nu);\theta,\sigma) \tag{C.6}$$

VG is a Levy process, like the Brownian motion, but unlike it, is a pure-jump process, due to the gamma time-change. Its characteristic function is:

$$\phi_{VG_t}(u) = E\left[e^{iuVG_t}\right] = \left(\frac{1}{1 - i\theta\nu u + (\sigma^2\nu/2)u^2}\right)^{t/\nu} \tag{C.7}$$

and its characteristic exponent is:

$$\psi_{VG(u)} = \frac{1}{\nu}\ln\left(1 - i\theta\nu u + (\sigma^2\nu/2)u^2\right) \tag{C.8}$$

At the second step, the VG process itself is time-changed. The time-change is independent of the VG process and is given by the integral of the mean-reverting CIR process.

The CIR process is solution to the following stochastic differential equation:

$$dy_t = k(\eta - y_t)dt + \lambda\sqrt{y_t}dW_t \tag{C.9}$$

Denote $Y_t = \displaystyle\int_0^t y_s ds$, then

$$VGSA(t;\sigma,\nu,\theta,k,\eta,\lambda) = VG(Y_t;\sigma,\nu,\theta) \tag{C.10}$$

The characteristic function of Y_t is:

$$\phi_{Y_t}(u) = E\left[e^{iuY_t}\right] = A(t,u)\exp B(t,u)y_0$$

where $A(t,u)$, $B(t,u)$ and γ are as given in Section 2.2. Since Y_t is independent of the VG process, the characteristic function of the VGSA process can be obtained via conditional expectation:

$$\phi_{VGSA_t}(u) = E\left[e^{iuVGSA_t}\right] = \phi_{Y_t}(i\psi_{VG(u)}) \tag{C.11}$$

Then the characteristic function of the log-price at horizon t under the VGSA process is:

$$E^Q\left[e^{iu\log(S_t)}\right] = e^{iu\log(S_0)+iurt}\frac{\phi_{VGSA_t}(u)}{\phi_{VGSA_t}(-i)} \tag{C.12}$$

Table 1. Fixed moneyness and maturity levels

Moneyness levels for five fixed maturites at which option returns are calculated as for each name. Moneyness is the ratio between option strike and spot.

Maturity	Puts			Calls		
1 m.	0.90	0.95	1	1	1.05	1.10
3 m.	0.85	0.90	1	1	1.10	1.15
6 m.	0.80	0.90	1	1	1.10	1.20
9 m.	0.80	0.90	1	1	1.10	1.20
12m.	0.75	0.85	1	1	1.15	1.25

Table 2. Option data

The table displays the names and ticker symbols of the options, used in the estimations. The third column shows average implied at-the-money volatility over 1997 - 2002 for each name. Names are later sorted according to implied volatility in forming portfolios of option returns. The last three columns show the proportion of three maturity groups in the average daily open interest for at- and out-of-the-money options over 1997 - 2002 for each name.

		Average	Maturity		
Company name	Ticker	implied vol.	< 2 m.	2-7 m.	> 7m.
Amgen	AMGN	0.43	0.37	0.50	0.14
American Express	AXP	0.37	0.37	0.48	0.14
AOL	AOL	0.56	0.35	0.49	0.16
Boeing	BA	0.34	0.24	0.52	0.24
Bank Index	BKX	0.30	0.70	0.29	0.01
Citibank	C	0.36	0.34	0.49	0.17
Cisco Systems	CSCO	0.52	0.38	0.49	0.13
Pharmaceutical Index	DRG	0.26	0.78	0.22	0.00
General Electric	GE	0.32	0.33	0.50	0.16
Hewlett-Packard	HWP	0.45	0.38	0.51	0.11
IBM	IBM	0.35	0.39	0.46	0.16
Intel	INTC	0.45	0.35	0.48	0.17
Lehman Brothers	LEH	0.48	0.35	0.53	0.12
Merryll Lynch	MER	0.43	0.35	0.49	0.17
Phillip Morris	MO	0.36	0.28	0.55	0.17
Merck	MRK	0.29	0.31	0.54	0.16
Microsoft	MSFT	0.39	0.30	0.48	0.22
National Semicond.	NSM	0.65	0.38	0.48	0.14
Nextel Communic.	NXTL	0.63	0.32	0.56	0.11
Oracle	ORCL	0.57	0.38	0.50	0.12
Pfizer	PFE	0.34	0.31	0.54	0.15
Russel 2000	RUT	0.24	0.68	0.32	0.01
S&P 500	SPX	0.23	0.35	0.47	0.18
Sun Microsystems	SUNW	0.55	0.37	0.48	0.15
Texas Instruments	TXN	0.53	0.31	0.45	0.24
Wal-mart Stores	WMT	0.35	0.38	0.48	0.14
Gold Index	XAU	0.45	0.57	0.40	0.04
Oil Index	XOI	0.24	0.84	0.16	0.00

Table 3. Estimation errors

The table displays, for each name, the average number of options per day, used in the estimations (after discarding days with less than 12 options), the proportion of estimations with average percentage error (A.P.E.) greater than 5% (also discarded), and the average A.P.E.. in the remaining days.

Ticker	Aver. daily options	Days with A.P.E.>5%	Remaining A.P.E.
AMGN	28	0.041	0.028
AXP	30	0.020	0.029
AOL	35	0.066	0.026
BA	26	0.042	0.033
BKX	161	0.015	0.024
C	30	0.044	0.032
CSCO	28	0.019	0.025
DRG	58	0.056	0.035
GE	35	0.040	0.031
HWP	27	0.039	0.030
IBM	42	0.019	0.025
INTC	36	0.009	0.025
LEH	26	0.045	0.027
MER	31	0.017	0.027
MO	33	0.079	0.035
MRK	28	0.050	0.030
MSFT	39	0.017	0.025
NSM	19	0.036	0.028
NXTL	23	0.031	0.025
ORCL	23	0.016	0.026
PFE	34	0.071	0.030
RUT	73	0.040	0.031
SPX	121	0.029	0.025
SUNW	34	0.007	0.024
TXN	31	0.004	0.025
WMT	30	0.080	0.030
XAU	35	0.021	0.033
XOI	53	0.092	0.035

Table 4. VIX and risk-neutral standard deviation

Panel A shows regression output for $\ln R_{t+30} = \alpha + \beta \ln VIX_t + \varepsilon_t$. R_{t+30} is the realized daily return volatility of S&P500 over a 30-day period starting at time t. Only non-overlapping intervals are involved. VIX_t is the CBOE's Volatility Index calculated at the beginning of each 30-day interval. t-statistics are in parentheses. The two sub-periods are 1/1/1997 - 6/30/2000 and 7/1/2000 - 12/31/2002. Panel B shows regression output for $\ln R_{t+1} = \alpha + \beta \ln SD_t + \varepsilon_t$. SD_t is risk-neutral standard deviation of S&P500 at 30-day horizon, calculated at the beginning of each 30-day interval.

Panel A. VIX

	1997 - 2002		1997 - 2000		2000 - 2002	
α	-0.52	(-1.17)	-0.62	(-0.99)	-0.52	(-0.76)
β	0.82	(4.94)	0.79	(3.42)	0.81	(3.12)
R^2	0.24		0.21		0.24	

Panel B. Risk-neutral standard-deviation

	1997 - 2002		1997 - 2000		2000 - 2002	
α	-0.54	(-1.23)	-0.62	(-0.97)	-0.57	(-0.86)
β	0.80	(4.92)	0.77	(3.33)	0.78	(3.14)
R^2	0.24		0.21		0.24	

Table 5. Average option excess returns

Panel A shows the average of expected excess returns to unhedged option across all names, in each of the strike and maturity groups. Daily returns are multiplied by 30 (monthly basis). E.g. -0.21 stands for -21% of the option price monthly. Each row refers to one of the five maturity groups. O-T-M columns refer to the most out-of-the-money puts / calls; A-T-M columns refer to at-the-money puts / calls. MID columns refer to puts / calls with intermediate moneyness (as in Table 1). Panel B shows the average of expected excess returns to delta-hedged options in the same moneyness and maturity groups. Daily returns are multiplied by 30 (monthly basis) and are now given in % of the spot price. E.g. -0.38 stands for -0.38% of spot monthly. The bottom part of each panel shows the respective t-statistics (average returns divided by standard deviation times the square root of the number of names). Averages are given for the entire 1997 - 2002 period.

Panel A. Average returns to unhedged options

	Puts			Calls		
Maturity	O-T-M	MID	A-T-M	A-T-M	MID	O-T-M
1 m.	-0.21	-0.16	-0.04	0.26	0.58	0.90
3 m.	-0.05	-0.05	-0.02	0.15	0.28	0.46
6 m.	-0.01	-0.01	-0.02	0.11	0.17	0.40
9 m.	0.02	-0.00	-0.01	0.10	0.16	0.38
12 m.	0.15	0.03	0.00	0.10	0.29	0.67

t-statistics (unhedged options)

Maturity	O-T-M	MID	A-T-M	A-T-M	MID	O-T-M
1 m.	-4.29	-3.70	-1.38	5.88	11.43	9.28
3 m.	-2.44	-2.86	-1.61	6.91	10.66	6.99
6 m.	-0.71	-0.80	-1.93	7.79	8.92	4.67
9 m.	1.20	-0.04	-1.28	8.36	7.42	4.00
12 m.	2.65	1.90	0.33	8.20	4.39	3.70

Panel B. Average returns to delta-hedged options

	Puts			Calls		
Maturity	O-T-M	MID	A-T-M	A-T-M	MID	O-T-M
1 m.	-0.38	-0.37	-0.16	-0.04	0.19	-0.01
3 m.	-0.09	-0.09	-0.06	0.04	0.09	0.06
6 m.	0.01	-0.01	-0.01	0.08	0.07	0.07
9 m.	0.07	0.08	0.07	0.13	0.12	0.14
12 m.	0.11	0.12	0.14	0.19	0.15	0.15

t-statistics (delta-hedged options)

Maturity	O-T-M	MID	A-T-M	A-T-M	MID	O-T-M
1 m.	-7.83	-5.25	-1.68	-0.48	2.48	-0.15
3 m.	-1.82	-1.49	-0.83	0.52	1.56	1.15
6 m.	0.18	-0.08	-0.16	1.47	1.29	1.56
9 m.	1.75	1.65	1.32	2.21	1.91	2.25
12 m.	2.37	2.16	2.20	2.81	2.17	2.35

Table 6. Volatility risk prices - one vs. two volatility factors
(all unhedged option returns time-series)

The table shows volatility risk prices estimated with two-step cross-sectional regressions on all 840 time-series of unhedged daily option returns for 1997 - 2002. The estimated relations are:

$$R^i = \alpha_i + \beta_i^M MKT + \beta_i^{1m} VOL^{1m} + \beta_i^L VOL^L + \varepsilon^i$$

$$E\left[R^i\right] = \lambda^M \beta_i^M + \lambda^{1m} \beta_i^{1m} + \lambda^L \beta_i^L + \gamma^i$$

At the second step regessions are run separately for each day and the estimates are then averaged. MKT denotes daily returns on S&P 500, VOL^{1m} denotes daily changes in one-month S&P 500 volatility and VOL^L denotes daily changes in one of the 3, 6, 9 or 12-month volatilities (i.e risk-neutral standard deviations). The λ-s are estimated risk prices for each of the risk factors. Shanken corrected t-statistics are shown for each risk price estimate. In parenthesis is the proportion of alphas in the first-pass regression, estimated to be significant at 5%. In square brackets is the adj. R^2 in regressing average returns on betas. The two panels show results for two and one volatility factors resp.

	Panel A. Two volatility factors			Panel B. One volatility factor	
	Risk price λ	t-stat.		Risk price λ	t-stat.
MKT	0.07	2.41	MKT	0.07	2.44
VOL 1m	-0.09	-1.99	VOL 1m	-0.03	-0.62
VOL 3m	0.09	1.83		(0.96)	[0.34]
	(0.17)	[0.39]			
			MKT	0.07	2.46
MKT	0.06	2.17	VOL 3m	0.05	1.04
VOL 1m	-0.12	-2.61		(0.96)	[0.35]
VOL 6m	0.17	3.20			
	(0.16)	[0.40]	MKT	0.07	2.34
			VOL 6m	0.12	2.28
MKT	0.06	1.99		(0.96)	[0.36]
VOL 1m	-0.14	-2.88			
VOL 9m	0.24	3.85	MKT	0.07	2.22
	(0.15)	[0.42]	VOL 9m	0.16	2.84
				(0.96)	[0.36]
MKT	0.05	1.90			
VOL 1m	-0.14	-2.90	MKT	0.06	2.14
VOL 12m	0.29	4.22	VOL 12m	0.21	3.26
	(0.16)	[0.42]		(0.95)	[0.38]

Table 7. Volatility risk prices - raw vs. orthogonal volatility risks
(all unhedged option returns time-series)

The table shows volatility risk prices estimated with two-step cross-sectional regressions on all 840 time-series of unhedged daily option returns for 1997 - 2002. The estimated relations are:

$$R^i = \alpha_i + \beta_i^M MKT + \beta_i^{1m} VOL^{1m} + \beta_i^L VOL^L + \varepsilon^i$$

$$E\left[R^i\right] = \lambda^M \beta_i^M + \lambda^{1m} \beta_i^{1m} + \lambda^L \beta_i^L + \gamma^i$$

At the second step regessions are run separately for each day and the estimates are then averaged. MKT denotes daily returns on S&P 500, VOL^{1m} denotes daily changes in one-month S&P 500 volatility and VOL^L denotes daily changes in one of the 3, 6, 9 or 12-month volatilities (volatilities here are risk-neutral standard deviations). The λ-s are estimated risk prices for each of the three risk factors. Shanken corrected t-statistics are shown for each risk price estimate. Panel A shows regressions with the volatility risk factors VOL^{1m} and VOL^L. Regressions in panel B use the component of VOL^{1m} orthogonal to MKT, and the component of VOL^L orthogonal to each of the other two factors.

	Panel A. Raw vol.		Panel B. Orthogonal vol.	
	Risk price λ	t-stat.	Risk price λ	t-stat.
MKT	0.07	2.41	0.07	2.42
VOL 1m	-0.09	-1.99	-0.06	-0.95
VOL 3m	0.09	1.83	0.35	3.74
MKT	0.06	2.17	0.06	2.18
VOL 1m	-0.12	-2.61	-0.10	-1.84
VOL 6m	0.17	3.20	0.39	4.57
MKT	0.06	1.99	0.06	2.00
VOL 1m	-0.14	-2.88	-0.13	-2.02
VOL 9m	0.24	3.85	0.42	4.88
MKT	0.05	1.90	0.05	1.92
VOL 1m	-0.14	-2.90	-0.13	-2.09
VOL 12m	0.29	4.22	0.43	4.92

Table 8. Average excess returns on delta-hedged option portfolios

Five portfolios are formed at each maturity by sorting names according to average implied volatility (see Table 1). Volatility quintiles are numbered from 1 (lowest volatility) to 5 (highest volatility). Average daily excess returns for 1997 - 2002 are multiplied by 30 (monthly basis) and given in %; e.g. -0.29 stands for -0.29% of spot monthly.

Maturity	1 Low vol.	2	3	4	5 High vol.
1 m	-0.29	-0.04	-0.03	0.00	0.02
3 m	-0.17	-0.03	-0.03	0.00	0.06
6 m	-0.10	-0.03	0.00	0.02	0.07
9 m	-0.07	-0.02	0.02	0.06	0.05
12 m	-0.03	-0.02	0.04	0.07	0.07

Table 9. Volatility risk prices - two volatility factors
(twenty five portfolios)

The table shows volatility risk prices estimated with GMM on twenty five portfolios of delta-hedged options
The moment conditions are:

$$g = \begin{bmatrix} E(R-\alpha-\beta^M MKT-\beta^{1m}VOL^{1m}-\beta^L VOL^L) \\ E\left[(R-\alpha-\beta^M MKT-\beta^{1m}VOL^{1m}-\beta^L VOL^L)MKT\right] \\ E\left[(R-\alpha-\beta^M MKT-\beta^{1m}VOL^{1m}-\beta^L VOL^L)VOL^{1m}\right] \\ E\left[(R-\alpha-\beta^M MKT-\beta^{1m}VOL^{1m}-\beta^L VOL^L)VOL^L\right] \\ E(R-\lambda^M\beta^M-\lambda^{1m}\beta^{1m}-\lambda^L\beta^L) \end{bmatrix} = 0$$

MKT denotes daily returns on S&P 500, VOL^{1m} denotes daily changes in one-month volatility, VOL^L denotes daily changes in one of 3, 6 9 or 12 month volatility (volatilities here are risk-neutral standard deviations). The λ-s are estimated risk prices for each of the three risk factors. z-statistics are distributed standard normal. Tilded factors (e.g. $\widetilde{VOL}$ 3m.) are the components of the respective raw factors, orthogonal to VOL^{1m}. p-values for the chi-squared test for pricing errors jointly equal to zero are in parenthesis. In square brackets is the adj. R^2 in regressing average returns on betas.

	1997 - 2002		1997 - 2000		2000 - 2002	
	Risk price λ	z-stat.	Risk price λ	z-stat	Risk price λ	z-stat
MKT	0.00	0.06	0.03	0.47	-0.01	-0.07
VOL 1m.	-0.27	-3.77	-0.33	-3.51	-0.11	-1.44
VOL 3m.	0.01	0.16	-0.02	-0.25	0.07	0.78
$\widetilde{VOL}$ 3m.	0.42	2.11	0.46	2.36	0.22	1.31
	(0.42)	[0.69]	(0.72)	[0.68]	(0.85)	[0.20]
MKT	0.02	0.32	0.08	1.22	-0.01	-0.15
VOL 1m.	-0.31	-4.07	-0.46	-3.67	-0.12	-1.60
VOL 6m.	0.12	1.57	0.19	1.40	0.09	1.39
$\widetilde{VOL}$ 6m.	0.46	3.15	0.68	2.79	0.26	1.96
	(0.85)	[0.86]	(0.97)	[0.91]	(0.90)	[0.23]
MKT	0.00	0.00	0.03	0.47	-0.02	-0.26
VOL 1m.	-0.29	-4.37	-0.43	-3.83	-0.12	-1.73
VOL 9m.	0.10	1.81	1.15	1.53	0.08	1.53
$\widetilde{VOL}$ 9m.	0.35	3.59	0.47	2.91	0.22	2.30
	(0.90)	[0.86]	(0.88)	[0.92]	(0.92)	[0.19]
MKT	-0.01	-0.21	0.01	0.12	-0.03	-0.34
VOL 1m.	-0.28	-4.43	-0.40	-3.88	-0.12	-1.81
VOL 12m.	0.11	2.01	0.16	1.79	0.08	1.46
$\widetilde{VOL}$ 12m.	0.30	3.53	0.39	2.92	0.19	2.40
	(0.85)	[0.85]	(0.80)	[0.90]	(0.92)	[0.16]

Table 10. Volatility risk prices - one volatility factor
(twenty five portfolios)

The table shows volatility risk prices estimated with GMM on twenty five portfolios of delta-hedged options. Themoment conditions are:

$$g = \begin{bmatrix} E(R-\alpha-\beta^M MKT-\beta^V VOL) \\ E\left[(R-\alpha-\beta^M MKT-\beta^V VOL)MKT\right] \\ E\left[(R-\alpha-\beta^M MKT-\beta^V VOL)VOL\right] \\ E(R-\lambda^M\beta^M-\lambda^V\beta^V) \end{bmatrix} = 0$$

MKT denotes daily returns on S&P 500, VOL denotes daily changes in volatility (volatilities are risk-neutral standard deviations at 1, 3, 6, 9 or 12-month horizons). The λ-s are estimated risk prices for each of the two risk factors. z-statistics are distributed standard normal. p-values for the chi-squared test for the pricing errors jointly equal to zero are in parenthesis. In square brackets is the adj. R^2 in regressing average returns on betas. The sub-periods are 1/1/1997 - 6/30/2000 and 7/1/2000 - 12/31/2002

	1997 - 2002		1997 - 2000		2000 - 2002	
	Risk price λ	z-stat.	Risk price λ	z-stat.	Risk price λ	z-stat.
MKT	-0.04	-0.72	-0.03	-0.39	-0.03	-0.43
VOL 1m.	-0.21	-4.31	-0.26	-4.09	-0.09	-1.54
	(0.24)	[0.63]	(0.30)	[0.65]	(0.88)	[0.10]
MKT	-0.06	-0.95	-0.06	-0.70	-0.01	-0.12
VOL 3m.	-0.20	-3.87	-0.26	-3.84	-0.03	-0.53
	(0.11)	[0.52]	(0.05)	[0.63]	(0.79)	[0.01]
MKT	-0.02	-0.35	0.01	0.16	0.01	0.11
VOL 6m.	-0.12	-2.23	-0.13	-1.87	0.01	0.19
	(0.00)	[0.42]	(0.05)	[0.57]	(0.63)	[0.01]
MKT	0.01	0.24	0.06	0.95	0.01	0.10
VOL 9m.	-0.06	-1.29	-0.05	-0.81	0.01	0.16
	(0.00)	[0.45]	(0.02)	[0.60]	(0.67)	[0.00]
MKT	0.02	0.51	0.07	1.24	0.00	0.06
VOL 12m.	-0.03	-0.75	-0.01	-0.25	0.00	0.02
	(0.00)	[0.47]	(0.02)	[0.61]	(0.78)	[0.00]

Table 11. Volatility risk prices - absolute market returns and one volatility factor (twenty five portfolios)

The table shows volatility risk prices estimated with GMM on twenty five portfolios of delta-hedged options. The moment conditions are analogous to those in Table 9. MKT denotes daily returns on S&P 500, $|MKT|$ is the absolute value of MKT, VOL denotes daily changes in one of 1, 3, 6, 9 or 12 month volatility The λ-s are estimated risk prices for each of the three risk factors. z-statistics are distributed standard normal. Tilded factors (e.g. $\widetilde{VOL}$ 1m.) are the components of the respective raw factors, orthogonal to $|MKT|$. p-values for the chi-square test for pricing errors jointly equal to zero are in parenthesis. In square brackets is the adj. R^2 in regressing average returns on betas. The two sub-periods are as in Table 9.

	1997 - 2002		1997 - 2000		2000 - 2002	
	Risk price λ	z-stat.	Risk price λ	z-stat.	Risk price λ	z-stat.
MKT	-0.08	-1.30	-0.03	-0.48	-0.05	-0.74
|MKT|	-0.16	-1.78	-0.03	-0.25	-0.16	-1.72
VOL 1m.	-0.20	-3.80	-0.26	-4.40	-0.04	-0.68
$\widetilde{VOL}$ 1m.	-0.32	-3.27	-0.38	-3.71	-0.07	-0.56
	(0.36)	[0.74]	(0.37)	[0.71]	(0.88)	[0.21]
MKT	-0.12	-1.68	-0.04	-0.53	-0.05	-0.58
|MKT|	-0.19	-2.25	0.04	0.39	-0.18	-1.94
VOL 3m.	-0.19	-3.16	-0.25	-4.27	0.01	0.12
$\widetilde{VOL}$ 3m.	-0.39	-2.84	-0.39	-3.35	-0.03	-0.18
	(0.31)	[0.69]	(0.07)	[0.66]	(0.83)	[0.21]
MKT	-0.09	-1.36	0.11	1.83	-0.04	-0.51
|MKT|	-0.20	-2.38	0.22	1.68	-0.17	-1.90
VOL 6m.	-0.13	-2.07	-0.09	-1.45	0.03	0.47
$\widetilde{VOL}$ 6m.	-0.26	-2.00	-0.04	-0.43	0.00	0.00
	(0.10)	[0.66]	(0.01)	[0.60]	(0.82)	[0.23]
MKT	-0.04	-0.74	0.16	2.80	-0.04	-0.54
|MKT|	-0.15	-1.98	0.28	1.99	-0.17	-1.89
VOL 9m.	-0.07	-1.28	-0.01	-0.20	0.03	0.55
$\widetilde{VOL}$ 9m.	-0.11	-1.28	0.07	0.89	0.01	0.04
	(0.04)	[0.65]	(0.01)	[0.61]	(0.85)	[0.23]
MKT	-0.02	-0.35	0.17	2.94	-0.04	-0.58
|MKT|	-0.13	-1.71	0.28	2.01	-0.17	-1.90
VOL 12m.	-0.04	-0.76	0.02	0.35	0.03	0.51
$\widetilde{VOL}$ 12m.	-0.05	-0.75	0.08	1.26	0.01	0.05
	(0.03)	[0.66]	(0.01)	[0.61]	(0.86)	[0.21]

Table 12. Average returns to calendar spreads

Panel A shows average returns to calendar spreads formed from short options with maturity at least 50 days and long options with the next available maturity and of the same type and strike. The positions are held for non-overlapping 30-day periods. The strikes of A-T-M options are within $\pm$ 5% of the spot at the beginning of of each 30-day period. O-T-M (I-T-M) options are at least 5% out-of-the-money (in-the-money) at the beginning of each 30-day period. Average spread returns for S&P500 alone and for all 28 names in the sample are shown. The sub-periods are 1/1/1997 - 6/30/2000 and 7/1/2000 - 12/31/2002. Panel B shows Sharpe ratios for calendar spreads in the same moneyness groups and periods.

Panel A. Average one-month returns to calendar spreads

	1997 - 2002		1997 - 2000		2000 - 2002	
	S&P500	All names	S&P500	All names	S&P500	All names
All puts	0.20	0.20	0.16	0.21	0.26	0.17
All calls	0.13	0.12	0.13	0.12	0.14	0.10
A-T-M puts	0.15	0.09	0.20	0.10	0.06	0.06
A-T-M calls	0.05	0.05	0.09	0.07	-0.03	0.03
O-T-M puts	0.16	0.15	0.02	0.13	0.33	0.17
O-T-M calls	0.27	0.20	0.37	0.25	0.15	0.14
I-T-M puts	0.56	0.38	0.70	0.46	0.41	0.26
I-T-M calls	0.13	0.09	0.03	0.07	0.27	0.11

Panel B. Sharpe ratios for calendar spreads

	1997 - 2002		1997 - 2000		2000 - 2002	
	S&P500	All names	S&P500	All names	S&P500	All names
All puts	0.41	0.34	0.37	0.36	0.46	0.37
All calls	0.25	0.31	0.26	0.35	0.24	0.29
A-T-M puts	0.49	0.29	0.64	0.32	0.21	0.24
A-T-M calls	0.20	0.24	0.44	0.33	-0.12	0.12
O-T-M puts	0.31	0.30	0.05	0.25	0.53	0.43
O-T-M calls	0.31	0.43	0.41	0.57	0.19	0.31
I-T-M puts	0.89	0.50	1.15	0.63	0.68	0.43
I-T-M calls	0.31	0.27	0.13	0.22	0.51	0.38